GaDei: On Scale-up Training As A Service For Deep Learning

Wei Zhang
weiz@us.ibm.com
IBM Research

Minwei Feng
mfeng@us.ibm.com
IBM Watson

Yunhui Zheng
zhengyu@us.ibm.com
IBM Research

Yufei Ren
yren@us.ibm.com
IBM Research

Yandong Wang
yandong@us.ibm.com
IBM Research

Ji Liu
ji.liu.uwisc@gmail.com
University of Rochester

Peng Liu
liup@us.ibm.com
IBM Research

Bing Xiang
bingxia@us.ibm.com
IBM Watson

Li Zhang
zhangli@us.ibm.com
IBM Research

Bowen Zhou
zhou@us.ibm.com
IBM Watson

arXiv:1611.06213v1 [stat.ML] 18 Nov 2016

Abstract

Deep learning (DL) training-as-a-service (TaaS) is an important emerging industrial workload. The unique challenge of TaaS is that it must satisfy a wide range of customers who have no experience and/or resources to tune DL hyper-parameters (e.g., mini-batch size and learning rate) , and meticulous tuning for each user's dataset is prohibitively expensive. Therefore, TaaS hyper-parameters must be fixed with values that are applicable to all users. IBM Watson Natural Language Classifier (NLC) service, the most popular IBM cognitive service used by thousands of enterprise-level clients around the globe, is a typical TaaS service. By evaluating the NLC workloads, we provide empirical evidence that only the conservative hyper-parameter setup (e.g., small mini-batch size and small learning rate) can guarantee acceptable model accuracy for a wide range of customers. We further justify theoretically why such a setup guarantees better model convergence in general. Unfortunately, the small mini-batch size causes a high volume of communication traffic in a parameter-server based system, which is the de facto DL scale-out/scale-up system design. We characterize the high communication bandwidth requirement of TaaS using representative industrial deep learning workloads and demonstrate that none of the state-of-the-art scale-up or scale-out solutions can satisfy such a requirement. We then present GaDei, an optimized shared-memory based scale-up parameter server design. We implemented three optimization strategies: (1) minimize memory copy by implementing a multiple-producer-single-consumer gradient queue in the shared-memory region mapped to the learner and parameter server processes, (2) in-place lock-free parameter server weights update, and (3) staging on device GPU memory buffers to enable asynchronous communication between parameter server and learners. We prove that the designed protocol is deadlock-free and it processes each gradient exactly once. Our implementation is evaluated on both commercial benchmarks and public benchmarks to demonstrate that it significantly outperforms the state-of-the-art parameter-server based implementation while maintaining the required accuracy and our implementation reaches near the best possible runtime performance, constrained only by the hardware limitation. Furthermore, to the best of our knowledge, GaDei is the only scale-up DL system that provides fault-tolerance.

[Copyright notice will appear here once 'preprint' option is removed.]

1. Introduction

When deployed on the cloud, deep learning (DL) training-as-a-service (TaaS) faces unique challenges: worldwide customers upload various types of training data and expect a model of high prediction accuracy returned in a short period of time. Unlike academic researchers, customers often have neither expertise nor resources to conduct time-consuming hyper-parameter (e.g., learning rate, mini-batch size) tuning. Hyper-parameter tuning itself is an unsolved and challenging research topic and the tuning process is usually prohibitively expensive [9, 52]. The goal of TaaS is not to provide each user a dedicated hyper-parameter setup and the companioning model, but to provide all users an unified DL model and the common hyper-parameter setup which still delivers the cutting-edge technology to the customers (in other words, no obvious accuracy loss). As a result, **industrial practitioners adopt conservative hyper-parameter setup** (e.g., small mini-batch size and small learning rate).

In a parameter server (PS) based DL system, such a conservative setup implies high-frequency communication with the PS. We provide a detailed analysis of the communication bandwidth requirement for real-world commercial workloads in Section 2.2 and Section 3. The analysis shows that the bandwidth requirement is beyond the capacity of advanced communication hardware (e.g., RDMA). Furthermore, with faster GPU devices and more efficient software library support, such as cuDNN [11], the communication cost of exchanging models between parameter servers and learners start to dominate the training cost, which renders scale-out deep learning training unattractive.

In addition, the staleness issue inherent in distributed deep learning systems makes scaling-out deep learning less cost-effective when training over a dozen GPU learners [39, 58]. One of the largest public-known commercial deep learning training system [5] uses 8 GPUs precisely for this reason. Consequently, **scale-up** deep learning training becomes a favored approach in the TaaS scenario.

Although powerful scale-up servers, such as NVIDIA DevBox, provides hardware platforms to improve the training performance, our evaluation (detailed in Section 6.3) reveals that the state-of-the-art scale-up software solutions are unable to make the best use of underlying hardware. Two representative open-source scale-up deep learning frameworks are mpiT[57], an MPI-based ASGD[1] framework; and DataParallelTable (DPT), a nccl-based (nccl[42] is a high-performance in-node GPU collectives implementation) SSGD[2] framework. The former solution incurs unnecessary memory copies between GPUs and MPI runtime and is unable to implement the lock-free update as proposed in HogWild![41]. The latter solution incurs synchronization barriers by forcing all GPUs to operate in lock-steps, which leads to the straggler problem. Consequently, the adoption of those mechanisms results in GPU stalls, which lead to severely degraded device utilization ($\leq 10\%$) on commercial workloads. Furthermore, none of the two solutions provide a fault-tolerance mechanism, which makes them undesirable for commercial adoption.

To solve these issues, we introduce GaDei, a high-performance scale-up parameter server design, aiming to efficiently coordinate the model synchronization among GPUs located on the same machine. GaDei strives to pipeline the entire model synchronization, overlapping all the model training and data movement phases to eliminate GPU stalls. Specifically, GaDei implements three system optimizations: (1) Communication via minimum memory copies (2) Lock-free Hogwild! style weights update rule (3) On-device double buffering, along with GPU multi-streaming, to pipeline model trainings and parameter movements. GaDei enables training with small mini-batch size, which mitigates the staleness issue to guarantee model convergence. By evaluating GaDei on a diverse set of real-world deep learning workloads, we demonstrate that GaDei is able to efficiently exploit the bandwidth offered by the commodity scale-up servers, providing faster convergence with significantly higher training speedups compared to existing open-source solutions, such as mpiT and DPT.

Overall, this work has made the following contributions:

1. We have identified the key challenges in designing a training-as-a-service system: Hyper-parameters must be set conservatively (e.g., small mini-batch size and high model communication frequency) to guarantee model accuracy.
2. We have designed and implemented GaDei, a highly-optimized parameter server system, to deliver scale-up and resilient training for TaaS workloads on multi-GPU servers. GaDei enables efficient multi-learner training for arbitrary type of neural networks (e.g., CNN, RNN). The design principle of GaDei is independent from the underlying gradient-calculation building-blocks and can complement any open-source DL frameworks (e.g., Torch[17], Caffe[29], and TensorFlow[4]).
3. We have proved that GaDei's system design guarantees both model convergence and deadlock-free. To the best of our knowledge, GaDei is the only scale-up parameter server design that provides both fault-tolerance and deadlock-free guarantee. We have systematically evaluated GaDei's performance by using 6 deep learning workloads with 3 state-of-the-art deep learning models. Evaluation results demonstrate GaDei often outperforms state-of-the-art solutions by an order of magnitude.

[1] Asynchronous Stochastic Gradient Descent (ASGD) is defined in Section 2.1

[2] Synchronous Stochastic Gradient Descent (SSGD) is defined in Section 2.1

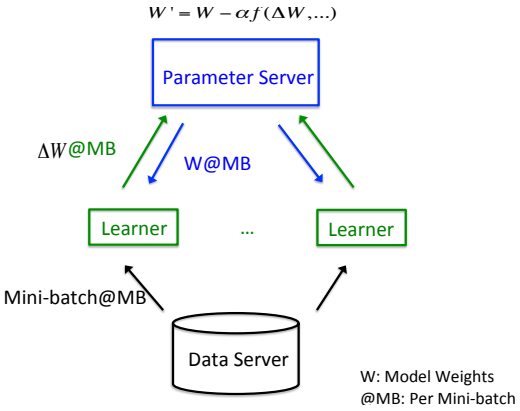

Figure 1: A typical parameter server architecture.

2. Background and Motivation

In this section, we introduce the background and define the terminologies (in bold font) used in this paper. Then we describe the characteristics of TaaS workloads. Finally, we theoretically justify why our design choice guarantees the acceptable model accuracy.

2.1 Terminology Definition

In essence, deep learning solves the following generic optimization problem

$$\min_{\theta} \quad F(\theta) := \frac{1}{N}\sum_{n=1}^{N} f_n(\theta),$$

where θ is the **parameter** (or **weights**) vector we are seeking, N is the number of samples, and $f_n(\theta)$[3] is the loss function for the n^{th} sample. Stochastic Gradient Descent(SGD) is the de facto algorithm to solve the deep learning optimization problem. SGD iterates over each training sample and applies Equation 1 to update weights. In Equation 1, i is the iteration number, k is the k-th parameter, ∇ is the differential operator, and α is the **learning rate**. Using a large learning rate may converge faster but it may also overshoot so that it does not converge at all, thus using a smaller learning rate is a safer choice in the production run. SGD passes through the entire training dataset several times until the model converges. Each pass is called an **epoch** (denoted as E). To improve computation efficiency, one can group a number of samples (i.e., a **mini-batch**, the size of one mini-batch is denoted as μ) and apply Equation 1 to update weights $\theta^{(k)}(i)$ for the $i+1$-th mini-batch.

$$\theta^{(k)}(i+1) = \theta^{(k)}(i) - \alpha\nabla\theta^{(k)}(i) \tag{1}$$

To accelerate deep learning training, practitioners usually adopt the **Parameter Server** (PS) architecture as illustrated in Figure 1. Each **learner** retrieves a mini-batch of training samples from data storage, calculates gradients and sends the gradients to the PS. PS then updates its weights using received gradients. Before calculating the next gradients, each learner pulls weights from the PS. We use λ to represent number of learners. Two most-widely adopted parameter server communication protocols are Synchronous SGD (**SSGD**) and Asynchronous SGD (**ASGD**). In SSGD, the PS collects gradients from each learner and then updates weights following the rule defined in Equation 2. SSGD is mathematically equivalent to SGD when $\lambda \times \mu$ is equal to the mini-batch size used in SGD.

$$\begin{aligned}\nabla\theta^{(k)}(i) &= \frac{1}{\lambda}\sum_{l=1}^{\lambda}\nabla\theta_l^{(k)} \\ \theta^{(k)}(i+1) &= \theta^{(k)}(i) - \alpha\nabla\theta^{(k)}(i)\end{aligned} \tag{2}$$

SSGD is not computationally efficient because the PS stalls the learners until it finishes collecting the gradients from all learners and updating the weights. ASGD relaxes such constraints by applying gradient update rule defined in Equation 3.

$$\begin{aligned}\nabla\theta^{(k)}(i) &= \nabla\theta_l^{(k)}, L_l \in L_1, ..., L_\lambda \\ \theta^{(k)}(i+1) &= \theta^{(k)}(i) - \alpha\nabla\theta^{(k)}(i)\end{aligned} \tag{3}$$

In ASGD, whenever PS receives a gradient from any learner, PS starts updating its weights. ASGD has the obvious runtime performance advantage over SSGD because learners do not wait for each other to start communicating with the PS. On the other hand, PS and the learners see different weights. PS always has the most up-to-date weights and the discrepancy between the weights used in a learner and the weights stored on PS is measured by staleness. When PS updates the weights, it increments the weights's (scalar) timestamp by 1; **staleness** is defined as the difference between the timestamp of the learner's weights and the timestamp of PS's weights. A large staleness can cause learners to mis-calculate the gradients, which leads to a slower convergence rate[19, 21, 53, 58]. We explain why training with a small mini-batch size can effectively reduce staleness in Section 2.3. Additionally, several recent works [24, 39, 57, 58] demonstrate that ASGD can converge to a similar model accuracy ($\pm 1\%$) as SGD after training with the same number of epochs, when the staleness is bounded in the system (typically up to a dozen of learners in the system). Assuming T_1 is the time for a single learner SGD algorithm to train E epochs and T_2 is the time for ASGD algorithm to train E epochs, ASGD **speed up** is defined as $\frac{T1}{T2}$. For a fair comparison, one must also certify that the model accuracy trained by ASGD is similar to that of SGD after E epochs.

2.2 Characteristics of the TaaS Workloads

In this section, we detail the characteristics of the training-as-a-service workloads by studying IBM Watson's natural language classification (NLC) service, which is the most popular service on IBM Watson's cognitive computing cloud and used by thousands of enterprise-level customers globally.

The NLC task is to classify input sentences into a target category in a predefined label set. NLC has been exten-

[3] $f(\theta)$ is typically in a form of a multi-layered neural network.

sively used in many practical applications, including sentiment analysis, topic classification, question classification, author profiling, intent classification, and even bug detection, and etc. State-of-the-art method for NLC is based on deep learning [25, 30, 32, 36, 48]. Figure 2 illustrates the deep learning model used in the NLC service.

NLC service is deployed as **"training-as-a-service"** in the cloud. After customers upload their in-house training data to the cloud, the NLC model training will be triggered in the background. The NLC model is ready to use after the training completes. Hence from the customer perspective, the turn around time is the model training time. Although the deep learning brings superior classification accuracy, one known issue is the time-consuming training phase, which can severely affect customer's experience. **Minimizing the training time has become IBM Watson NLC's top priority.**

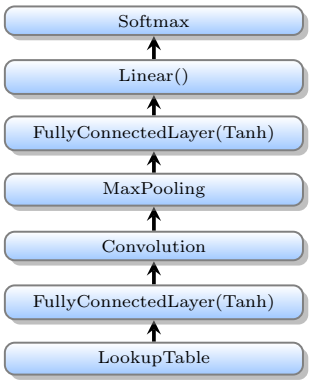

Figure 2: Convolutional Neural Network (CNN) based deep learning model for natural language classification tasks.

In order to improve the performance for neural network training, previous work has attempted to use scale-out frameworks to coordinate learners distributed on different computing nodes. Good speedup on large-scale tasks like ImageNet[22] has been reported. However, this type of tasks does not represent the commercial TaaS workloads. By examining the real TaaS workloads, we find that corpus size of training sets are generally less than 10k for most use-cases and usually they come with a diverse set of labels. The reason is that in practice annotating training data is expensive for most customers. In addition, we have identified the following characteristics that are critical to the quality of trained model:

(1) Large batch sizes can incur significant accuracy loss: It is well-known that using large mini-batch can improve GPU utilization and incur less demand for communication; however, this runtime performance improvement must not sacrifice model accuracy. Our field study reveals that the deep learning method used in this paper is on average 3%-6% more accurate than other much less computation-intensive methods (e.g., SVM) for NLC tasks. Therefore, an accuracy loss larger than 3%-6% will invalidate the use of deep learning for NLC tasks. To study the impact of different batch sizes on model accuracy, we use 4 representative NLC datasets (the detailed description of each task is given in Table 5) and evaluate the accuracy under different batch sizes in Figure 3. The batch size is increased from 1 to 128. Experimental results show that using large batch sizes would result in unacceptable accuracy loss for three out of four cases. When using large mini-batch size, models for challenging workloads (a large amount of labels and little training data for each label) such as Joule and Watt do not even converge. Other researchers have also observed that large batch size slows down convergence [31, 38, 55].

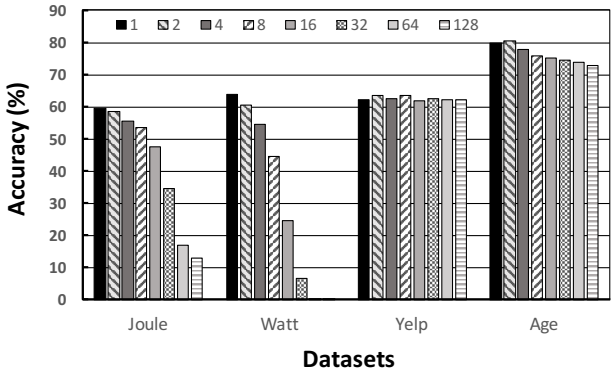

Figure 3: Accuracy performance of different datasets under different batch sizes. Mini-batch size ranges from 1 to 128. The number of training epochs is fixed at 200. Using large mini-batch size severely decreases model accuracy.

(2) Low communication frequency decreases model quality: Communication bandwidth is typically much lower in the cloud than in the HPC systems. One way of mitigating the communication bottleneck is to allow each learner to process many mini-batches before it synchronizes with the parameter server. However, although less-frequent communication can efficiently increase the GPU utilization, but it can severely decrease model accuracy as shown in Table 1. Intuitively, less frequent communication causes higher discrepancy (i.e., staleness) between learners and the PS and it will decrease model accuracy. Therefore, learners should communicate with the PS as frequently as possible. Ideally, each learner shall communicate with the PS after each mini-batch training.

(3) Conservative hyper-parameter configuration is imperative: Previous research focuses on how to speedup training for one specific dataset and heavy hyper parameter tuning is required to achieve the best possible accuracy. For example, Table 2 shows the typical hyper-parameter setups for training CIFAR and ImageNet with AlexNet[34] model. The setups vary greatly for different datasets. In contrast, the TaaS users have neither expertise nor resources for hyper-parameter tuning. Customers just upload the training data and then expect a well trained model to be ready in a short period of time. As a result, the hyper-parameters have to be preset to fixed values to cover diverse use cases. Ta-

Workload	Accuracy (%)		GPU utilization (%)	
	CI=1	CI=8	CI=1	CI=8
Joule	57.70	39.80	7.90	35.80
Watt	57.80	55.10	4.90	29.40
Age	74.10	63.60	4.00	24.90
MR	77.20	50.00	5.50	33.70
Yelp	56.90	20.00	8.80	42.90

Table 1: Accuracy loss and GPU utilization with different communication interval in the parameter server based training framework. CI is the communication interval, measured as the number of mini-batches each learner has processed before it communicates with the parameter server. High CI can improve GPU utilization but severely affects model accuracy. This experiment is conducted using the mpiT package, on a 12 NVIDIA K20-GPU cluster with a 10Gb/s interconnect.

ble 3 records the hyper-parameter used in NLC. For thousands of different datasets, NLC adopts a much simpler and more conservative setup. Note this is a significant difference from commonly evaluated workloads (e.g., CIFAR and ImageNet), where hyper-parameter tuning is specific to each dataset/model and usually is a result of a multi-person-years effort[9, 52]. In TaaS, **conservative configurations with small batch size, high communication frequency, and small learning rate are commonly adopted to satisfy a wide range of users.**

Dataset	CIFAR		ImageNet	
Mini-batch size		128		256
Number of epochs		100		40
Learning Rate(LR)	Epoch	LR	Epoch	LR
	E1-E25	1	E1-E18	0.01
	E26-E50	0.5	E19-E30	0.005
	E51-E75	0.25	E31-E40	0.001
	E75-E100	0.125		

Table 2: The typical hyper-parameter setup when training CIFAR and ImageNet with AlexNet model. To achieve best possible accuracy, researchers conduct heavy dataset-specific hyper-parameter tuning (e.g., sophisticated learning rate adaption schemes).

Dataset size	small	medium	large
Mini-batch size	2	4	32
Number of epochs		200	
Learning rate		0.01	

Table 3: The hyper-parameter setup in NLC. Instead of heavy tuning for each dataset as shown in Table 2, NLC uses a simpler and more conservative setup (e.g., small mini-batch size, small learning rate and large number of epochs) to satisfy all the users. Users are categorized in 3 groups based on the number of their training samples: small ($<$ 10K), medium (10K – 100K), and large ($>$ 100K).

In summary, the above workload characteristics study suggests that adopting the scale-out solutions such like distributed parameter server based frameworks is unsuitable for a set of industry deep learning tasks that require small mini-batch size and high frequency of model exchange.

2.3 Theoretical Justification of Using Small Mini-batch

In the previous section, we have empirically demonstrated training with large batch size can cause unacceptable accuracy loss. Based on a recent theoretical study [39], we now justify that why training with small batch size can counter system staleness and is desired for distributed deep learning in general.

Theorem 1. *[39] Under certain commonly used assumptions, if the learning rate is chosen in the optimal way and the staleness σ is bounded by*

$$\sigma \leq O\left(\sqrt{E/\mu^2}\right) \tag{4}$$

where E is the total number of epochs and μ is the mini-batch size, then the asynchronous parallel SGD algorithm converges in the rate

$$O(\sqrt{1/E}).$$

The result suggests the tolerance of the staleness T relies on the mini-batch size μ. First note that the convergence rate $O(\sqrt{1/E})$ is optimal. The prerequisite is that the staleness σ is bounded by $O(\sqrt{E/\mu^2})$. The staleness σ is usually propositional to the total number of learners. To satisfy the condition in (4), either μ should be small enough or the epoch number K should be large enough. In other words, given the number of learners and the total epoch number (or the total computational complexity), small mini-batch size is preferred. In addition, it also explains why small mini-batch size is potentially preferred even for SGD (running on a single worker), since SGD with mini-batch size μ can be considered as running Async-SGD with mini-batch size 1 with μ workers (it implies that the staleness is $O(\mu)$). Thus, SGD with large mini-batch size is equivalent to Async-SGD with a small mini-batch size but a large staleness.

3. Communication Bandwidth Requirement in TaaS

In this section, we measure the computation time for different values of μ over a wide range of NLC workloads and public image recognition workloads. We then calculate the minimum memory bandwidth requirement to achieve any speedup when ASGD protocol is employed.

Each learner's execution loop consists of three components: T_{train} (gradient calculation), T_{pull} (pull weights), T_{push}(push gradients). Each parameter server's execution loop contains three components: $T_{receive}$ (receive gradients), T_{apply} (apply weights update), and T_{send} (send weights). When μ is large, $T_{train} + T_{pull} + T_{push} \gg T_{receive} + T_{apply} + T_{send}$; when μ is small, time spent on PS becomes the critical path. In ASGD, sending weights and receiving gradients operations may overlap. To achieve any speedup,

we then must have $T_{train} \geq (T_{receive} + T_{apply})$ [4]. Note that the apply update operation is memory-bound level 1 BLAS operation. Combined memory bandwidth between GPU and CPU (gradients transfer) and memory bandwidth used in CPU DRAM (weights update) are of the same order of magnitude. Further, gradients and weights are of the same size. We now can infer the required overall communication bandwidth to observe any speedup is at least $\frac{2 \times ModelSize}{TrainTime_{perminibatch}}$. For NLC workload and image recognition workload, Table 4 records Training time Per Epoch (TPE), Training Samples number (N), Model Size; and calculates the minimum Required Bandwidth (RB) to observe any speedup. As we can see, it requires at least 10GB/s bandwidth to achieve any speedup for NLC workloads with the appropriate mini-batch size. In addition, to achieve X-fold ($X > 1$) speedup, we need to multiply RB by a factor of X. Such a demanding bandwidth is beyond the capacity of advanced network techniques (e.g., RDMA). Note RB is also quite close to the peak memory bandwidth (PCI-e / DRAM), which indicates any extra memory copy may make speedup impossible. Thus, it is natural to infer that the only viable PS architecture is a tightly coupled multi-GPU system collocated on the same server that minimizes data copies and enables learners to asynchronously push gradients and pull weights.

Why existing scale-up solutions are insufficient ? Among popular open-source deep learning frameworks, Caffe[29], Torch[17] and TensorFlow[4] support multi-GPU training on the same node. However, they are designed for tasks where heavy hyper-parameter tuning is allowed so a larger mini-batch size may be appropriate (e.g., 256). It is easy to see from Table 4 that it requires much higher communication bandwidth to support a small mini-batch than to support a large mini-batch. In addition, Caffe and TensorFlow only support SSGD on one node. We have demonstrated in Section 2.2 that some of the workloads require mini-batch size to be as small as 2, which means Caffe and TensorFlow can at most make use of 2 GPUs (e.g., each GPU works with a mini-batch size of 1). Torch is the only open-source DL framework that supports both SSGD (via DPT) and ASGD (via mpiT) on a single-node. However, as demonstrated in Section 6.3, neither DPT nor mpiT can efficiently use the memory bandwidth on the same node. Furthermore, none of the existing solutions provide a fault-tolerance mechanism in the scale-up setting.

4. Design and Implementation

4.1 Overall Design

GaDei strives to minimize memory copy and enable high-concurrency to maximize communication bandwidth utilization. Figure 4 depicts its design. To minimize memory copy, PS and learners use a shared-memory region to exchange gradients and weights. Each learner has a fixed number of slots in the producer-consumer queue, thus the entire system can be viewed as λ[5] single-producer-single-consumer queues. PS updates weights in place (i.e., HogWild! style). We use 4 `openmp` threads and unroll weights update loop 8 times to maximize DRAM throughput. To maximize system concurrency, each learner creates two additional threads – push thread and pull thread. On the same GPU device, each learner maintains an on-device gradient staging buffer so that after a learner finishes gradient calculation it can store the gradients in the buffer and continue the next gradient calculation without waiting for the completion of push. On-device memory bandwidth is usually several hundreds of GB/s, which is much faster than device-host memory bandwidth (typically $\sim$10 GB/s). By buffering gradients on the same device, learner can train continuously while the push thread is pushing gradients to PS. Similarly, each learner also maintains a weights staging buffer on the same device. Learners do not communicate with each other, they only communicate with parameter server.

Figure 5 details the necessary logic of PS and learners. We use the following naming conventions: variables that start with 'g_' (e.g., g_env, g_param_ptr) represent shared variables between PS and Learners, other variables (e.g., pullCnt, pushCnt) are shared variables between learner's main thread and its communication threads. PS (Figure 5a) iterates over the gradient queues in a round-robin fashion and it busy-loops when all the queues are empty. PS does not yield CPU via a conditional variable wait, because PS demands the most CPU cycles to process gradients and it is beneficial to have PS takeover gradients whenever they are ready. Figure 5c illustrates the logic of learner main thread. It calculates gradients on GPU's default stream. The learner main thread communicates with push thread (Figure 5d) via a producer-consumer queue (Line 105 - 107 in Figure 5c and corresponding Line 203-206, 220-222 in Figure 5d) of size 1, i.e. variable *pushCnt* in Figure 5c and Figure 5d alternates between 0 and 1. The push thread operates on a separate stream so that it can send gradients in concurrent with the learner thread calculating the gradients. Similarly the learner main thread communicates with the pull thread (Figure 5b) via a producer-consumer-queue (Line115-125 in Figure 5c and in Figure 5b) of size 1, i.e., variable *pullCnt* in Figure 5c and Figure 5b alternates between 0 and 1. If the weights in the learner thread is current (i.e. has the same timestamp as the weights on PS), pullThd skips the pull request in this iteration. By default, the PS updates weights at Line 7 in Figure 5a in a lock-free fashion (e.g., an incarnation of HogWild! algorithm[41]). GaDei also supports protecting weight updates from concurrent pulling via a read-write-lock.

[4] (i)Apply update and receive gradients cannot overlap, since apply update can only start when gradients are fully received (ii) Assuming learner can push gradients and receive weights instantaneously

[5] λ is the number of Learners.

μ	Joule		Watt		Age		Yelp		CIFAR		ImageNet	
	2.46K,7.69MB		7K,20.72MB		68.48K,72.86MB		500K,98.60MB		50K,59,97MB		1280K,244.48MB	
	TPE (sec)	RB (GB/s)	TPE (sec)	RB (GB/s)	TPE (sec)	RB (GB/s)	TPE (sec)	RB (GB/s)	TPE (sec)	RB (GB/s)	TPE (sec)	RB (GB/s)
1	5.61	7.00	25.22	11.50	574.80	17.36	4376.12	22.53	N/A*	N/A*	N/A*	N/A*
2	3.22	6.10	14.11	10.28	309.51	16.12	2367.04	20.83	792.45	3.78	26957.54	11.61
4	1.77	5.54	7.46	9.72	169.99	14.68	1299.42	18.97	502.63	2.98	15596.41	10.03
8	1.00	4.89	4.06	8.93	104.39	11.95	802.04	15.37	356.46	2.10	9499.38	8.24
16	0.75	3.27	2.79	6.49	77.79	8.02	571.10	10.79	290.95	1.29	7294.30	5.36
32	0.72	1.69	2.38	3.82	69.52	4.49	451.61	6.82	261.23	0.72	5769.76	3.39
64	0.80	0.76	2.37	1.91	77.14	2.02	403.69	3.82	245.68	0.38	5014.85	1.95
128	1.00	0.31	2.77	0.82	100.22	0.78	402.97	1.91	234.38	0.20	4784.22	1.02

Table 4: Minimum bandwidth requirement to see any speedup. μ is mini-batch size, TPE stands for Time Per Epoch, RB stands for Required Bandwidth. *Both CIFAR and ImageNet use models that use Batch Normalization (BN), which requires $\mu \geq 2$.

Note that $cudaStreamSynchronize()$ invoked during the execution is to make sure the memory was flushed in place (either between host and device or within the same device) w.r.t the same stream. As a result, the corresponding memory copy between CPU/GPU or within GPU strictly follows programming order, which is necessary for our protocol verification, and which will be described in the next section.

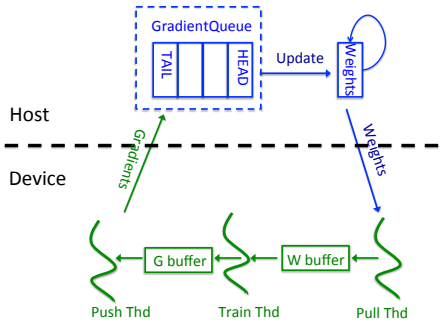

Figure 4: Overview of GaDei's Design

4.2 Verification of GaDei's communication protocol

It is difficult to detect, avoid and fix concurrency bugs[40]. GaDei relies on heavy communication between CPU threads, CPU-GPU interaction, and multi-stream operation within the same GPU device. It is imperative to verify the correctness of its communication protocol. In this section, we prove GaDei is deadlock-free in Theorem 2 and verify GaDei's liveness property in Theorem 3.

Lemma 1. *The value of pushCnt and pullCnt can only be 0 or 1. For the idx-th learner, $0 \leq$ g_env[idx].cnt $\leq$ g_env[idx].sz.*

Proof. We will show that pushCnt and pullCnt cannot be larger than 1 or smaller than 0. Line 110 (Fig. 5c) is the only place where pushCnt can be incremented. If pushCnt can be larger than 1 then there must be an iteration where pushCnt is 1 before executing line 110, since the increment is 1 per iteration. However, this is impossible because line 110 is not reachable due to the condition check loop at line 105. Similarly, line 220 (Fig. 5d) is the only place where pushCnt can be decremented. It cannot be smaller than 0 due to the loop at line 204. Therefore, pushCnt can only be 0 or 1. The claims about pullCnt and g_env[idx].cnt can be proved in the same way. □

Lemma 2. *Once signaled, a thread blocked by a condition wait (line 106, 117, 205, 209 or 305 in Fig. 5) will wake up and exit the corresponding condition check loop.*

Proof. We will show the loop conditions do not hold when signals are sent. For example, only line 222 (Fig. 5d) can wake up the condition wait on pushEmpty at line 106 . According to Lemma 1, pushCnt can only be 0 or 1. When sending the signal, pushCnt is always 0 (due to line 220) and thus invalidates the loop condition at line 105. Therefore, the training thread will exit the condition check loop. Similarly, we can prove the claim is true for other condition waits. □

Lemma 3. *The wait-for graph formed by condition waits only, i.e. pthread_cond_wait, is acyclic and thus deadlock-free.*

Proof. The directed edges in Fig. 6 represent the wait-for relations among wait and signal statements. The edges in blue and green form two cycles. We will show some edges in a cycle cannot exist at the same time.

In the blue cycle, the edge from line 205-2 to 111 represents that the push thread waits until pushCnt != 0. The blue edge from 106-2 to 221 indicates the training thread waits until pushCnt != 1. Given that pushCnt can only be 0 or 1 (Lemma 1), the above two conditions cannot be true at the same time. Therefore, these two edges cannot exist together and the cycle in blue is infeasible. In the green cycle, the edge from 305-2 to 124 illustrates the pull thread waits until pullCnt = 0. The edge from 117-2 to 317 indicates the training thread waits until pullCnt = 1. Similarly, this cycle is infeasible too. □

```
1.  /** Parameter Server Thread **/
2.  while (true) {
3.    if (g_env[idx].cnt == 0) {
4.      idx = (idx + 1) % workerNum;
5.    } else {
6.      float* src = g_buf[idx][g_env[idx].use_ptr];
7.      update(g_param_ptr, src, ...);
8.      g_param_version++;
9.      (g_env[idx]).use_ptr =
10.         ((g_env[idx]).use_ptr + 1) % (g_env[idx].sz);
11.     pthread_mutex_lock(&((g_env[idx]).gradMtx));
12.     g_env[idx].cnt--;
13.     pthread_cond_signal(&(g_env[idx].gradEmpty));
14.     pthread_mutex_unlock(&(g_env[idx].gradMtx));
15.     idx = (idx + 1) % workerNum;
16.   }
17. }
```

(a) Host (Parameter Server Thread)

```
301. /** Pull Thread **/
302. while(true){
303.   pthread_mutex_lock(&pullMtx);
304.   while(pullCnt == 1){
305.     pthread_cond_wait(&pullEmpty, &pullMtx);
306.   }
307.   if(p_version >= g_param_version){
308.     ...
309.   } else {
310.     if(p_version < g_param_version){
311.       cudaMemcpy(..., g_param_ptr, ...);
312.       cudasyncStream(...);
313.       p_version = g_param_version;
314.     }
315.   }
316.   pullCnt++;
317.   pthread_cond_signal(&pullFill);
318.   pthread_mutex_unlock(&pullMtx);
319. }
```

(b) Pull Thread

```
101. /** Training Thread **/
102. while(true){
103.   T *grad = calGradeints();
104.   pthread_mutex_lock(&pushMtx);
105.   while(pushCnt == 1){
106.     pthread_cond_wait(&pushEmpty, &pushMtx);
107.   }
108.   cudaMemcpyDevice2Device(grad_buf_gpu, grad, ...);
109.   cudaStreamSynchronizeDefaultStream();
110.   pushCnt++;
111.   pthread_cond_signal(&pushFill);
112.   pthread_mutex_unlock(&pushMtx);
113.
114.   T *weights = getWeightsStorage();
115.   pthread_mutex_lock(&pullMtx);
116.   while(pullCnt == 0){
117.     pthread_cond_wait(&pullFill, &pullMtx);
118.   }
119.   if(recvNewWeightsFlag){
120.     cudaMemcpyDevice2Device(weights, ...);
121.     cudaStreamSynchronizeDefaultStream();
122.   }
123.   pullCnt--;
124.   pthread_cond_signal(&pullEmpty);
125.   pthread_mutex_unlock(&pullMtx);
126. }
```

(c) Training Thread

```
201. /** Push Thread **/
202. while(true){
203.   pthread_mutex_lock(&pushMtx);
204.   while(pushCnt == 0){
205.     pthread_cond_wait(&pushFill, &pushMtx);
206.   }
207.   pthread_mutex_lock(&(g_env[tid].gradMtx));
208.   while ((g_env[tid].cnt) == (g_env[tid].sz)){
209.     pthread_cond_wait(&(g_env[tid].gradEmpty),
210.         &(g_env[tid].gradMtx));
211.   }
212.   pthread_mutex_unlock(&(g_env[tid].gradMtx));
213.   cudaMemcpy(g_buf[tid][g_env[tid].fill_ptr], ...);
214.   cudasyncStream(...);
215.   (g_env[tid]).fill_ptr =
216.       ((g_env[tid]).fill_ptr + 1) % g_env[tid].sz;
217.   pthread_mutex_lock(&(g_env[tid].gradMtx));
218.   g_env[tid].cnt++;
219.   pthread_mutex_unlock(&(g_env[tid].gradMtx));
220.   pushCnt--;
221.   pthread_cond_signal(&pushEmpty);
222.   pthread_mutex_unlock(&pushMtx);
223. }
```

(d) Push Thread

Figure 5: Details of the parameter server, learner main thread (training thread), push thread and pull thread.

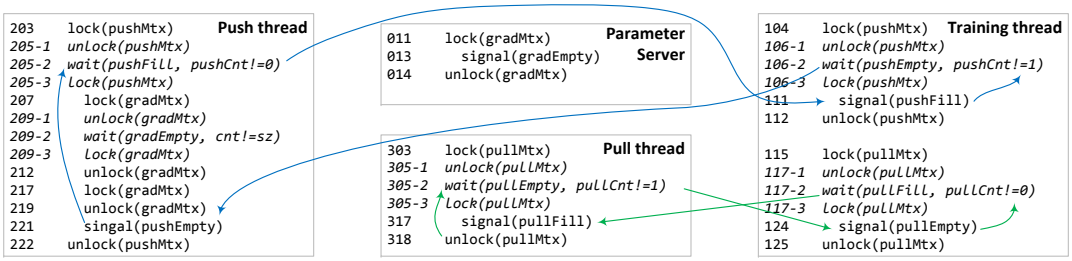

Figure 6: Synchronization operations extracted from Fig. 5. The line numbers refer to the same lines in Fig. 5. Lines in shape of i-1, i-2 and i-3 are equivalently transformed from the condition check loop at line i in Fig. 5. Operation *wait(s, c)* means the thread is blocked until it's waken up by signal *s* and condition *c* is true.

Lemma 4. *Mutex lock operations in Fig. 5 are deadlock-free.*

Proof. Since the condition check loops (lines 105-107, 116-118, 204-206, 208-211 and 304-306 in Fig. 5) do not have side effects, we equivalently rewrite the condition check loops and summarize synchronization operations in Fig. 6.

Now consider operations based on mutex locks in Fig. 6. One necessary conditions for deadlocks is *hold-and-wait* [16]. If threads are not holding one resource while waiting for another, there is no deadlock. However, only the push thread can be in a hold-and-wait state (lines 207, 209-3 and 217). Therefore, mutex lock operations cannot introduce deadlocks. □

Theorem 2. *GaDei is dealock-free.*

Proof. Lemma 3 and 4 show that neither condition waits nor mutex lock operations can cause deadlocks. Now we consider them together. In Fig. 6, condition waits at 106-2,

205-2, 305-2 and 117-2 are not in any atomic region and thus are isolated from mutex locks. Hence, they cannot introduce deadlocks. The only remaining case is the condition wait at 209-2. It is guarded by lock pushMtx and waits for the signal sent at 013. However, line 13 is protected by a different lock gradMtx. This doesn not satisfy the hold-and-wait condition. Therefore, it cannot lead to deadlocks either. □

Theorem 3. *The parameter server thread processes each gradient exactly once.*

Proof. The push thread of learner idx shares gradients with the parameter server using a share array g_env[idx]. Essentially, it is an array-based FIFO queue, where g_env[idx].cnt indicates the number of gradients in queue. It is straightforward to see *neither the push thread nor the parameter server can access the same memory location in two consecutive iterations*.

Similar to the proof for Lemma 1, we can show that $0 \leq g_env[idx].cnt \leq g_env[idx].sz$ so that $g_env[idx].sz$ is the max size of the queue. In addition, as the parameter server only processes a learner's gradients if its queue is not empty (lines 6-14), the read pointer ($g_env[idx].use_ptr$) can never be ahead of the write pointer ($g_env[tid].fill_ptr$). So, *the parameter server cannot read an outdated gradient and use it more than once*. Similarly, when the queue reaches its max size, the push thread must wait (lines 208-211). So, *the push thread cannot overwrite unprocessed gradients in the queue*. □

4.3 Fault-tolerance

In GaDei, each learner and the PS has its own address space, learners and PS communicate via mmap-ed memory. This approach is similar to Grace[8], which transforms multi-threaded program to multi-process program communicating via shared memory. When GaDei starts, learners and PS mmap the same memory file, which pre-allocated gradient queues, shared weights, and thread related synchronization variables (e.g., mutexes and condition variables, both with PTHREAD_PROCESS_SHARED attributes set).

PS periodically communicates with the watchdog process to log its progress and checkpoint the parameters. When n learners unexpectedly die ($n < \lambda$, λ is the number of learners), PS continues to process gradients collected from alive learners so that failures from the dead learners are naturally isolated. Note PS is stateless in that it only needs to process a fixed number of gradients without considering which learners produced the gradients.

If all learners die, PS no longer makes progress, thus the watchdog process kills the PS and restarts PS and learners from the last checkpoint. If a learner dies when holding a lock, PS will hang when it tries to grab the lock. Thus the watchdog process will later detect the failure and take action. Alternatively, one may set the robust attribute of pthread mutex so that when PS is grabbing the lock it can notice the failed learner and skip checking that learners' gradient queue.

4.4 Additional Design Decisions

Computation optimization inside GaDei. The major computation in GaDei's PS server is element-wise vector addition for updating the global weights. We unroll the weights update loop on the parameter server side 8 times in the OpenMP parallel section, and found that it outperforms the non-unrolled version by 30%. In addition, it outperforms the version in which we directly use the Streaming SIMD Extension (SSE) instructions.

Half-precision floating point operations. Recent research work[27, 28] demonstrate that it is possible to train DNNs with lower precision and still obtain comparable model accuracy. Comparing to 32-bit single precision, the 16-bit half-precision data format will decrease parameter server processing burden. The latest GPU architecture may support half-precision 16-bit float point operation, while general purpose CPUs do not. We integrated software-based 16-bit CPU floating point operations in GaDei. However, It slowed down the computation by a factor of 5, which makes any speedup impossible for our workloads.

Lock-free producer-consumer queue. A lock-free producer-consumer queue was considered in our design. Lock-free design can minimize the inter-process/thread latency, but would consume a CPU core in 100% for each learner. As a result, the computation threads have to compete with the lock-free communication threads, and this will impede the computation performance significantly. In GaDei, when learners have to wait for the server to process gradients, yielding CPUs via conditional variable wait can save CPU cycles for PS computation threads to process gradients.

5. Methodology

5.1 Software

We use open source toolkit Torch [17] as the building block to calculate the gradients of neural nets. Torch is a scientific computing framework based on Lua/LuaJIT with both CPU and CUDA backends. Torch has been widely used both in industry (e.g., Facebook, Google Deepmind, and IBM) and in academic community. Researchers at Google[4], Bosch[6], and Facebook[14] have benchmarked several commonly used open-source deep learning frameworks (e.g. Caffe, Theano, Torch, and Tensorflow) and found that Torch usually outperforms other frameworks. In addition, Torch is the only framework that supports ASGD on one-node. **Thus to enable further speedup on top of Torch presents a bigger challenge for the system design.**

5.2 Benchmark

We use four representative NLC datasets in our evaluation. Two datasets Joule and Watt are in-house customer

	Data Size	Label Size	Type	Description	Model Size	Network
Joule	2.4k	311	NLC	Question answering task in insurance domain; label represents answers.	7.69MB	CNN
Watt	7k	3595	NLC	Question answering task in online service domain; label represents answers.	20.72MB	CNN
Yelp	500k	5	NLC	Customer review classification; label represents the star the customer assigns to the business.	98.60MB	CNN
Age	68k	5	NLC	Author profiling task; the label represents the age range of the author.	72.86MB	CNN
MR	8.6k	2	NLC	Sentiment analysis of movie reviews; label represents positive/negative attitude of the audience.	14.27MB	CNN
CIFAR	50k	10	IR	Classify images into predefined categories.	59.97MB	VGG
ImageNet	1280k	1000	IR	Classify images into predefined categories.	244.48MB	AlexNet

Table 5: Task description and corpus statistics. NLC: natural language classification; IR: image recognition. CNN is the convolutional neural networks depicted in Figure 2. AlexNet [34] and VGG [50] are the de facto standard models in the field of image recognition.

datasets. The other two datasets Age[1] and Yelp[2] are publicly available datasets. The Joule and Watt datasets represent the typical small datasets present in NLC workloads. Age dataset represents the medium-size dataset. Yelp dataset represents the large-size dataset. The learning rate is fixed at 0.01 across workloads, GaDei trains each NLC workload for 200 epochs. The neural network model used for the four datasets is presented in Figure 2. To demonstrate that GaDei can be used as a drop-in replacement for tools that solve non-TaaS tasks, we also conducted experiments on image recognition tasks CIFAR [33] and ImageNet [46]. We train CIFAR using VGG model and we train ImageNet using AlexNet model. We use the widely-adopted hyper-parameter setup, as described in [13, 56], to train CIFAR and ImageNet tasks.Table 5 records the task description and training data statistics.

5.3 Hardware

The experiments have been conducted on the Softlayer cloud[6]. The server is equipped with two Intel Xeon E5-2690-V3 processors. Each processor has 12 real cores, clocked at 2.66GHz per core. To enable the best possible PS CPU processing speed, we turn off SMT. The CPU memory capacity is 128GB, with peak memory bandwidth 40GB/s. There are two NVIDIA Tesla K80s installed on the server. Each K80 has two GPUs. Totally there are four GPUs on the server with a total 16 TFlops. The bus interface of K80 is PCIe 3.0 x16, with a 12Gbps bi-directional bandwidth each lane.

6. Experiment Results

In this section, we evaluate GaDei's ability to achieve good model accuracy and runtime performance, against widely used state-of-the-art open-source multi-GPU implementation DPT and mpiT on both commercial workloads and public workloads. Section 6.1 demonstrates how the model accuracy progresses w.r.t training epoch. Section 6.2 illustrates the runtime performance of GaDei. Section 6.3 illustrates how GaDei achieves good speedup on challenging NLC workload, while other state-of-the-art tools cannot achieve any speedup. GaDei can achieve speedup using much smaller mini-batch size than any other tools. Section 6.4 discusses GaDei's ability to handle fault-tolerance and GPU over-subscription.

[6] http://www.softlayer.com

6.1 Convergence result

In Figure 7 we plot the model accuracy w.r.t the training epochs when using 1,2,3,4 learners. Joule and Watt converge to 60% , Age 80%, Yelp 62%, CIFAR 90%, ImageNet 55%. Model accuracy reaches the same level ($\pm 1\%$) of accuracy as the single learner system within the same number of training epochs.

Summary GaDei converges to the same level of accuracy as the single-learner SGD using the same number of epochs. This demonstrates a tightly-coupled system such as GaDei, when using small minibatch size, can mostly avoid the staleness issue introduced in a typical ASGD system.

6.2 Speedup Results and Memory Bandwidth Analysis

Workload	Model Size (MB)	Epochs	Running Time (hrs)
Joule	7.69	200	0.18
Watt	20.72	200	0.78
Age	72.86	200	9.44
Yelp	98.60	200	24.99
Cifar	59.97	100	6.53
ImageNet	244.48	40	52.56

Table 6: Single-learner performance baseline on one GPU.

Table 6 records the single-learner performance baseline for comparison. In Figure 8 the speedup performance of GaDei is plotted. When running on challenging commercial IBM Watson NLC workloads, GaDei can achieve on average

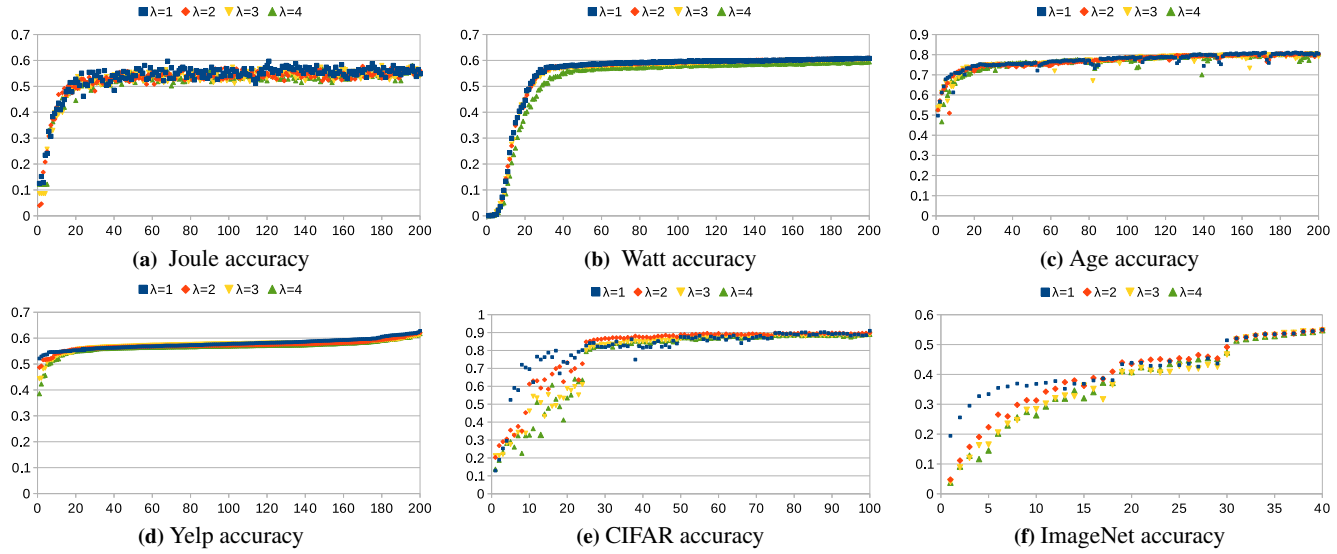

(a) Joule accuracy (b) Watt accuracy (c) Age accuracy

(d) Yelp accuracy (e) CIFAR accuracy (f) ImageNet accuracy

Figure 7: Model accuracy w.r.t the training epochs when using 1,2,3,4 GPUs. Joule and Watt converge to 60% , Yelp 62%, Age 80%, CIFAR 90%, ImageNet 55%. GaDei does not introduce accuracy loss which satisfies our expectation.

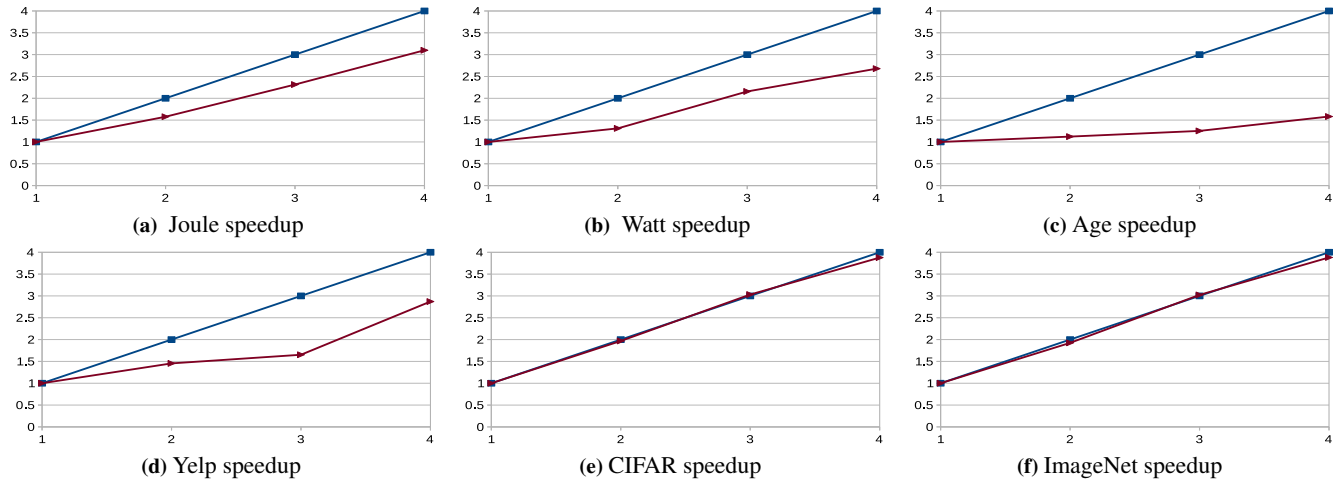

(a) Joule speedup (b) Watt speedup (c) Age speedup

(d) Yelp speedup (e) CIFAR speedup (f) ImageNet speedup

Figure 8: Blue line is oracle linear speedup. Red line is the speedup performance of GaDei .

1.5X - 3X speedup when using up to 4 learners. When running on public image recognition benchmark tasks, GaDei nearly achieves linear speedup. Dividing the total amount of data transferred between learners and parameter server by the total runtime, the memory bandwidth utilized by GaDei is reported in Table 7. When running Watson workloads on 4 GPUs, GaDei sustains 36-55GB/s bandwidth, which is close to the hardware limit.

Summary GaDei achieves linear speedup on public dataset and model, and achieves good speedup on challenging commercial workload. GaDei comes close to saturating the hardware memory bandwidth.

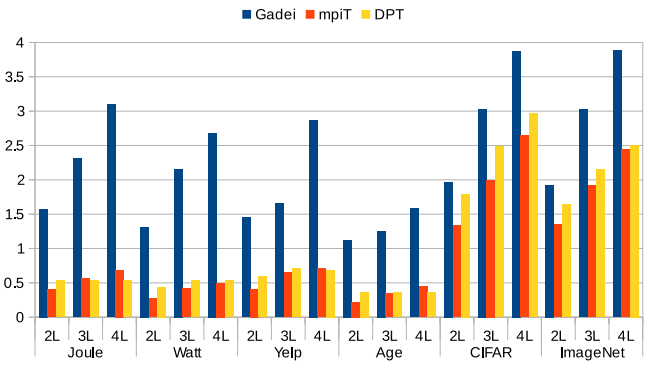

Figure 9: Speedup comparison among GaDei, mpiT and DPT. Y-axis is the speedup, X-axis represents different workload under $\lambda = 2, 3, 4$ settings. GaDei consistently outperforms other tools.

Workload	Memory Bandwidth Utilized(GB/s)		
	$\lambda = 2$	$\lambda = 3$	$\lambda = 4$
Joule	18.56	27.25	36.48
Watt	26.93	44.33	55.07
Age	32.94	36.76	46.42
Yelp	19.93	22.64	39.34
CIFAR	0.78	1.21	1.45
ImageNet	3.98	6.25	8.02

Table 7: Memory bandwidth utilization. λ is the number of learners.

Workload	Gadei	mpiT	DPT
Joule	1	16	32
Watt	1	16	64
Age	1	32	32
Yelp	16	64	128
CIFAR	2	2	8
ImageNet	2	8	8

Table 8: Minimum mini-batch size used by a system so that it can converge to expected accuracy while still observing speedup.

6.3 Compare with DPT and mpiT

We compare GaDei's speedup with that of mpiT and DPT, two state-of-the-art scale-up deep learning training frameworks. Figure 9 shows the comparison. GaDei consistently outperforms other tools. For NLC workload, DPT and mpiT actually slow down the execution. DPT has inferior performance because it is a SSGD style implementation, PS blocks all learners when updating weights, whereas GaDei implements ASGD style synchronization protocol. mpiT has inferior performance because it does not minimize memory copies (i.e., there is at least one extra copy from PS/learners to MPI runtime), and its message-passing style implementation makes HogWild! lock-free weights update impossible. Table 8 is the stress test: we decrease the batch size until thre is no longer a speedup. Recall as described in Section 2.3, the smaller batch size a system can support, the higher probability a deep learning model can reach a desirable model accuracy. GaDei supports much smaller batch size, often the smallest possible size that is constrained only by the underlying model (except for the Yelp dataset), than other tools.

Summary GaDei significantly outperforms state-of-the-art open source implementations. It achieves good speedup on commercial workloads whereas existing open source implementations slow down the execution. On public benchmark, GaDei also outperforms existing open-source alternatives. GaDei can speedup workload using batch size that is an order of magnitude smaller than other state-of-the-art tools.

6.4 Fault tolerance and over-subscription of learners

We randomly kill learners and verify GaDei can always finish training with desired model accuracy. In contrast, mpiT runs on top of MPI [3], which traditionally does not provide fault-tolerance mechanism. DPT orchestrates multiple learners in the same process, thus when one learner fails, the entire process is killed. Further, nccl implements blocking collective operations, and their behavior in the presence of failure is not defined.

GaDei supports over-subscription of GPUs , i.e., run λ learners over N GPUs, where $\lambda > N$. Recurrent neural networks, such as long short-term memory (LSTM), are known to be difficult to efficiently parallelize due to complicated computation dependencies. When such a model is adopted, GPUs usually operate at much lower efficiency compared to running CNNs. By oversubscribing GPUs, GaDei can fully utilize GPU resources. In another task, where LSTM is utilized, GaDei achieves 7-fold speedup when running 8 learners on 4 GPUs.

Summary GaDei supports fault-tolerance and GPU over-subscription. To the best of our knowledge, GaDei is the only scale-up deep learning system that supports fault-tolerance.

7. Related work

Deep learning has seen tremendous success in image recognition [35], natural language processing [23], speech translation [15] and gaming [49]. The parameter server [51] based approach is the de facto method to scale out such training tasks in a distributed environment. Several research works optimize PS performance and tackle fault-tolerance problems [12, 21, 37] in a CPU-only enviroment. As GPU-based deep learning frameworks [7, 17, 29] offer a much better cost-effective solution, GPU-based scale-out PS architecture, such as Mariana [59], is optimized for distributed GPU environment. GeePS [20] overcomes the memory capacity limit on the GPU device by loading a part of a model (instead of the entire model) that is necessary for computation to GPU at a given point, and treats CPU memory as a large data-storage.

Different from previously published work, IBM Watson Cognitive Service receives very different types of training data from worldwide customers instead of providing service using a centralized model. Consequently, it is imperative to set conservative hyper-parameter configuration. This requires extremely high communication bandwidth that renders scale-out solutions infeasible. State-of-the-art scale-up solutions, such as nccl-based SSGD implementation DPT[18] and MPI-based ASGD implementation mpiT[57] incur relatively large communication overhead due to intermediate data copy. Our solution solves these issues by min-

imizing the data copy and making the whole system highly concurrent.

Big code[10, 43] is an emerging research area, where researchers apply machine learning , and especially deep learning, to mine large code repository. Big code has seen great success in different areas of programing language and programming analysis, such like code completion[26, 44], type inference[45], and code similarity inspection[47, 54]. GaDei is a system-based approach to accelerating deep learning training for any type of the neural network (e.g., CNN, RNN), which is used by big code researchers.

8. Conclusion

In this paper, we focus on the system design challenges for emerging training-as-a-service (TaaS) workloads. By analyzing the characteristics of representative industrial workloads, we identify that to satisfy diverse customer requirements, a TaaS system needs to choose conservative hyper-parameter setup (e.g., small mini-batch size). We provide both empirical evidence and theoretical justification for such a design choice. We then characterize the communication bandwidth requirement for TaaS workloads and conclude that none of the state-of-the-art solutions can satisfy this requirement.

We present GaDei, a scale-up deep learning framework, that maximizes communication bandwidth utilization in a tightly-coupled system. GaDei enables efficient multi-learner training for arbitrary type of neural networks (e.g., CNN, RNN). We further verify the correctness of the GaDei's communication protocol. Our evaluation results demonstrate that GaDei significantly outperforms state-of-the-art scale-up solutions on industrial workloads and public workloads, usually by an order of magnitude. In addition, GaDei provides fault-tolerance, which is missing in other scale-up solutions.

References

[1] Author profiling. http://pan.webis.de/clef16/pan16-web/author-profiling.html.

[2] Yelp challenge. https://www.yelp.com/dataset_challenge.

[3] Mpi 3.0 standard. www.mpi-forum.org/docs/mpi-3.0/mpi30-report.pdf, 2012.

[4] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, GA, 2016. USENIX Association. ISBN 978-1-931971-33-1. URL `https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi`.

[5] D. Amodei, R. Anubhai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, J. Chen, M. Chrzanowski, A. Coates, G. Diamos, E. Elsen, J. Engel, L. Fan, C. Fougner, A. Y. Hannun, B. Jun, T. Han, P. LeGresley, X. Li, L. Lin, S. Narang, A. Y. Ng, S. Ozair, R. Prenger, S. Qian, J. Raiman, S. Satheesh, D. Seetapun, S. Sengupta, C. Wang, Y. Wang, Z. Wang, B. Xiao, Y. Xie, D. Yogatama, J. Zhan, and Z. Zhu. Deep speech 2 : End-to-end speech recognition in english and mandarin. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 173–182, 2016. URL `http://jmlr.org/proceedings/papers/v48/amodei16.html`.

[6] S. Bahrampour, N. Ramakrishnan, L. Schott, and M. Shah. Comparative study of caffe, neon, theano, and torch for deep learning. *CoRR*, abs/1511.06435, 2015. URL `http://arxiv.org/abs/1511.06435`.

[7] F. Bastien, P. Lamblin, R. Pascanu, J. Bergstra, I. J. Goodfellow, A. Bergeron, N. Bouchard, and Y. Bengio. Theano: new features and speed improvements. Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop, 2012.

[8] E. D. Berger, T. Yang, T. Liu, and G. Novark. Grace: Safe multithreaded programming for c/c++. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '09, pages 81–96, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-766-0. doi: 10.1145/1640089.1640096. URL `http://doi.acm.org/10.1145/1640089.1640096`.

[9] J. Bergstra and Y. Bengio. Random search for hyper-parameter optimization. *J. Mach. Learn. Res.*, 13:281–305, Feb. 2012. ISSN 1532-4435. URL `http://dl.acm.org/citation.cfm?id=2188385.2188395`.

[10] P. Bielik, V. Raychev, and M. T. Vechev. Programming with "big code": Lessons, techniques and applications. In T. Ball, R. Bodk, S. Krishnamurthi, B. S. Lerner, and G. Morrisett, editors, *1st Summit on Advances in Programming Languages, SNAPL 2015, May 3-6, 2015, Asilomar, California, USA*, volume 32 of *LIPIcs*, pages 41–50. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. ISBN 978-3-939897-80-4. doi: http://dx.doi.org/10.4230/LIPIcs.SNAPL.2015.41.

[11] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer. cudnn: Efficient primitives for deep learning. *CoRR*, abs/1410.0759, 2014. URL `http://arxiv.org/abs/1410.0759`.

[12] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman. Project Adam: Building an efficient and scalable deep learning training system. OSDI'14, pages 571–582, 2014. ISBN 978-1-931971-16-4.

[13] S. Chintala. *Multi-GPU ImageNet Training*. URL `https://github.com/soumith/imagenet-multiGPU.torch`.

[14] S. Chintala. *convnet-benchmark*, 2016. URL `github.com/soumith/convnet-benchmarks`.

[15] K. Cho, B. Van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, 2014. Association for Computational Linguistics. URL `http://www.aclweb.org/anthology/D14-1179`.

[16] E. G. Coffman, M. Elphick, and A. Shoshani. System deadlocks. *ACM Comput. Surv.*, 3(2):67–78, June 1971. ISSN 0360-0300. doi: 10.1145/356586.356588. URL `http://doi.acm.org/10.1145/356586.356588`.

[17] R. Collobert, K. Kavukcuoglu, and C. Farabet. Torch7: A matlab-like environment for machine learning. In *BigLearn, NIPS Workshop*, 2011.

[18] T. community. *Torch Data Parallel Table*. URL `https://github.com/torch/cunn/blob/master/DataParallelTable.lua`. Available at `https://github.com/torch/cunn/blob/master/DataParallelTable.lua`.

[19] H. Cui, J. Cipar, Q. Ho, J. K. Kim, S. Lee, A. Kumar, J. Wei, W. Dai, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and E. P. Xing. Exploiting bounded staleness to speed up big data analytics. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference*, USENIX ATC'14, pages 37–48, Berkeley, CA, USA, 2014. USENIX Association. ISBN 978-1-931971-10-2. URL `http://dl.acm.org/citation.cfm?id=2643634.2643639`.

[20] H. Cui, H. Zhang, G. R. Ganger, P. B. Gibbons, and E. P. Xing. Geeps: scalable deep learning on distributed gpus with a gpu-specialized parameter server. In *Proceedings of the Eleventh European Conference on Computer Systems, EuroSys 2016, London, United Kingdom, April 18-21, 2016*, page 4, 2016. doi: 10.1145/2901318.2901323. URL `http://doi.acm.org/10.1145/2901318.2901323`.

[21] J. Dean, G. S. Corrado, R. Monga, K. Chen, M. Devin, Q. V. Le, M. Z. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Y. Ng. Large scale distributed deep networks. In *NIPS*, 2012.

[22] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.

[23] M. Feng, B. Xiang, M. R. Glass, L. Wang, and B. Zhou. Applying deep learning to answer selection: A study and an open

task. In *Proceedings of the 2015 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU 2015)*., volume abs/1508.01585, 2015. URL http://arxiv.org/abs/1508.01585.

[24] M. Feng, B. Xiang, and B. Zhou. Distributed deep learning for question answering. In *The 25th ACM International Conference on Information and Knowledge Management (CIKM 2016)*, Indianapolis, IN, USA, October 2016. Association for Computing Machinery.

[25] I. Goodfellow, Y. Bengio, and A. Courville. Deep learning. Book in preparation for MIT Press, 2016. URL http://www.deeplearningbook.org.

[26] X. Gu, H. Zhang, D. Zhang, and S. Kim. Deep api learning. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2016, pages 631–642, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4218-6. doi: 10.1145/2950290.2950334. URL http://doi.acm.org/10.1145/2950290.2950334.

[27] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep learning with limited numerical precision. In D. Blei and F. Bach, editors, *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1737–1746. JMLR Workshop and Conference Proceedings, 2015. URL http://jmlr.org/proceedings/papers/v37/gupta15.pdf.

[28] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *5th International Conference on Learning Representations*, 2016.

[29] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22Nd ACM International Conference on Multimedia*, MM '14, pages 675–678, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-3063-3. doi: 10.1145/2647868.2654889. URL http://doi.acm.org/10.1145/2647868.2654889.

[30] N. Kalchbrenner, E. Grefenstette, and P. Blunsom. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 655–665, Baltimore, Maryland, June 2014. Association for Computational Linguistics. URL http://www.aclweb.org/anthology/P14-1062.

[31] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *CoRR*, abs/1609.04836, 2016. URL http://arxiv.org/abs/1609.04836.

[32] Y. Kim. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar, October 2014. Association for Computational Linguistics. URL http://www.aclweb.org/anthology/D14-1181.

[33] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

[34] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. 2012.

[35] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[36] Y. Lecun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 5 2015. ISSN 0028-0836. doi: 10.1038/nature14539.

[37] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su. Scaling distributed machine learning with the parameter server. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pages 583–598, Broomfield, CO, Oct. 2014. USENIX Association. ISBN 978-1-931971-16-4. URL https://www.usenix.org/conference/osdi14/technical-sessions/presentation/li_mu.

[38] M. Li, T. Zhang, Y. Chen, and A. J. Smola. Efficient mini-batch training for stochastic optimization. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14, pages 661–670, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2956-9. doi: 10.1145/2623330.2623612. URL http://doi.acm.org/10.1145/2623330.2623612.

[39] X. Lian, Y. Huang, Y. Li, and J. Liu. Asynchronous Parallel Stochastic Gradient for Nonconvex Optimization. *Neural Information Processing Systems*, June 2015.

[40] S. Lu, S. Park, E. Seo, and Y. Zhou. Learning from mistakes: A comprehensive study on real world concurrency bug characteristics. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS XIII, pages 329–339, New York, NY, USA, 2008. ACM. ISBN 978-1-59593-958-6. doi: 10.1145/1346281.1346323. URL http://doi.acm.org/10.1145/1346281.1346323.

[41] F. Niu, B. Recht, C. R, and S. J. Wright. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Neural Information Processing Systems*, 2011.

[42] Nvidia. *NCCL: Optimized primitives for collective multi-GPU communication*. URL https://github.com/NVIDIA/nccl.

[43] V. Raychev. Learning from large codebases. In *ETH Zrich, PhD thesis*, 2016. doi: http://dx.doi.org/10.3929/ethz-a-010737712.

[44] V. Raychev, M. Vechev, and E. Yahav. Code completion with statistical language models. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '14, pages 419–428, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2784-8. doi: 10.1145/2594291.2594321. URL http://doi.acm.org/10.1145/2594291.2594321.

[45] V. Raychev, M. Vechev, and A. Krause. Predicting program properties from "big code". In *Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '15, pages 111–124, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3300-9. doi:

10.1145/2676726.2677009. URL http://doi.acm.org/10.1145/2676726.2677009.

[46] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.

[47] H. Sajnani, V. Saini, J. Svajlenko, C. K. Roy, and C. V. Lopes. Sourcerercc: Scaling code clone detection to big-code. In *Proceedings of the 38th International Conference on Software Engineering*, ICSE '16, pages 1157–1168, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3900-1. doi: 10.1145/2884781.2884877. URL http://doi.acm.org/10.1145/2884781.2884877.

[48] J. Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015. doi: 10.1016/j.neunet.2014.09.003. Published online 2014; based on TR arXiv:1404.7828 [cs.NE].

[49] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–503, 2016. URL http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html.

[50] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2015.

[51] A. Smola and S. Narayanamurthy. An architecture for parallel topic models. *Proc. VLDB Endow.*, 3(1-2):703–710, 2010. ISSN 2150-8097. doi: 10.14778/1920841.1920931.

[52] J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2951–2959. Curran Associates, Inc., 2012. URL http://papers.nips.cc/paper/4522-practical-bayesian-optimization-of-machine-learning-algorithms.pdf.

[53] F. W. Suyog Gupta, Wei Zhang. Model accuracy and runtime tradeoff in distributed deep learning: A systematic study. In *IEEE International Conference on Data Mining*, 2016.

[54] M. White, M. Tufano, C. Vendome, and D. Poshyvanyk. Deep learning code fragments for code clone detection. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*, ASE 2016, pages 87–98, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3845-5. doi: 10.1145/2970276.2970326. URL http://doi.acm.org/10.1145/2970276.2970326.

[55] D. R. Wilson and T. R. Martinez. The general inefficiency of batch training for gradient descent learning. *Neural Netw.*, 16(10):1429–1451, Dec. 2003. ISSN 0893-6080. doi: 10.1016/S0893-6080(03)00138-2. URL http://dx.doi.org/10.1016/S0893-6080(03)00138-2.

[56] S. Zagoruyko. *VGG CIFAR*. URL https://github.com/szagoruyko/cifar.torch.

[57] S. Zhang, A. E. Choromanska, and Y. LeCun. Deep learning with elastic averaging sgd. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 685–693. Curran Associates, Inc., 2015. URL http://papers.nips.cc/paper/5761-deep-learning-with-elastic-averaging-sgd.pdf.

[58] W. Zhang, S. Gupta, X. Lian, and J. Liu. Staleness-aware async-sgd for distributed deep learning. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 2350–2356, 2016.

[59] Y. Zou, X. Jin, Y. Li, Z. Guo, E. Wang, and B. Xiao. Mariana: Tencent deep learning platform and its applications. *Proc. VLDB Endow.*, 7(13):1772–1777, 2014. ISSN 2150-8097. doi: 10.14778/2733004.2733082.